The ASC Roadrunner System at LANL: Near-term capacity with a future hybrid accelerated option

Manuel Vigil, Roadrunner Project Manager, HPC-DO

Los Alamos National Laboratory (LANL), in partnership with International Business Machines (IBM), is in the early stages of developing and deploying the Roadrunner High Performance Computing system. LANL Divisions involved include High Performance Computing (HPC); Computer, Computational, and Statistical Sciences (CCS); Computing, Telecommunications, and Networking (CTN); Applied Physics (X); and Theoretical (T).

Fig. 1.
The Roadrunner base system has been installed at LANL.

The primary goals for the Roadrunner system are as follows:

- Providing a large "capacity-mode" computing resource for LANL weapons simulations
- Implement and option for a petascale hybrid accelerated architecture capable of supporting the future LANL workload.
- Becoming a lead participant in the industry-wide path toward hybrid accelerated computing devices for HPC.

The Roadrunner Project has three phases. During Phase 1, already underway, LANL has acquired, installed, and deployed more than 81 teraFLOPS of a base capacity system to provide capacity computing cycles in the near term. Phase 2 is a technology refresh and assessment of the final system. Phase 3 consists of an option to install a hybrid computing architecture that has the potential for significant improvements in the price/performance curve to help meet future Advanced Simulation and Computing (ASC) computing requirements.

Roadrunner Base System

The Roadrunner base system was delivered in September 2006. LANL completed acceptance testing of the system in December 2006. The Roadrunner base system has 14 connected units (~71 TeraFLOPS) for the classified computing environment and 2 connected units (~10 TeraFLOPS) for the unclassified computing environment. Each connected unit is 144 Opteron X64 processors from Advanced Micro Devices (AMD) connected with a high-speed InfiniBand 4X interconnect fabric. The system in the unclassified computing environment also includes initial testbeds of the IBM Cell Broadband Engine™ (Cell BE) processing elements. These testbeds are being used for initial applications and systems software work in advance of the optional hybrid system in Phase 3.

Following acceptance, the system has been undergoing a focused system integration effort for assimilation into the LANL classified computing environment and initial applications testing. This integration includes infrastructure planning and deployment of key network and I/O capabilities to accommodate Roadrunner in the Metropolis Center for Modeling and Simulation, and installing and testing the production software stack (compilers, debuggers, Message Passing Interface, I/O libraries, resource manager, etc.) required by LANL applications. System monitoring and performance tools and processes for tracking system reliability and usage statistics are also being integrated.

The system was transitioned to the classified computing environment following security accreditation in May 2007. Several key weapons applications are running on the base system in anticipation of meeting a Level-2 ASC Milestone by June 2007. The system will formally transition to production status in the summer of 2007. Roadrunner will more than double the enduring capacity computing available for nuclear weapons applications at LANL.

Hybrid Computing Architecture – A New era in Scientific Simulation

Phase 2 is planning and assessing the hybrid architecture system targeted for deployment in Phase 3. The advanced hybrid architecture system will be both Opteron processors and the IBM Cell Broadband Engine™ (Cell BE) processing elements. Roadrunner is the first supercomputer system to use this hybrid processor architecture.

Los Alamos and IBM have been working together to improve the architecture of the Phase 3 system by improving performance both in the interconnect and between the hybrid nodes. The final classified system, planned for 2008, will have Opteron and IBM Cell Blades directly connected for sustained petaFLOP performance.

The final hybrid system will demonstrate a new paradigm for HPC, one in which accelerators are used to increase application per-formance rather than just adding more processors. Los Alamos and IBM are jointly developing a hybrid system programming model that will first be used in Roadrunner but will form the basis for future hybrid systems as well. Early initial performance on the Cell processors show great potential.

For more information contact Manuel Vigil at mbv@lanl.gov.

Funding Acknowledgements

This research was supported by the NNSA tri-Lab Advanced Simulation and Computing Program.

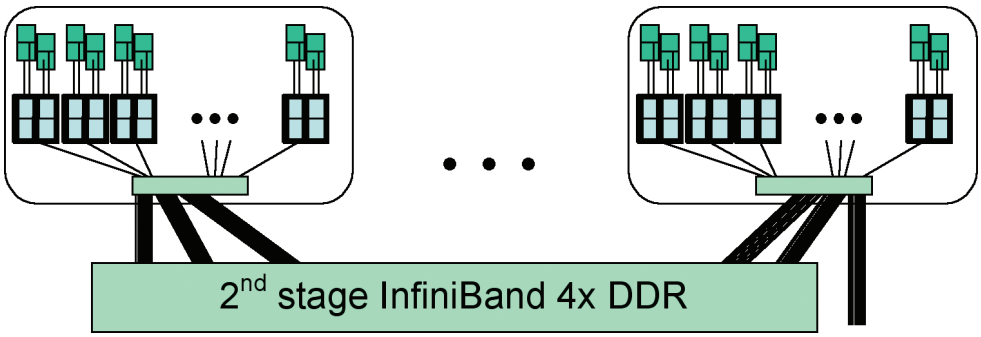

Fig. 2.
The Roadrunner hybrid architecture with Cell Blade accelerators.